

Dell EMC

D-MSS-DS-23 Exam

Dell Midrange Storage Solutions Design 2023

**Questions & Answers
Demo**

Version: 4.0

Question: 1

What is done during the preconfigure step of a PowerStore block import migration?

- A. Source system is added to the PowerStore.
- B. Destination volume is created.
- C. Network connectivity is set up.
- D. Import schedule is set.

Answer: C

Explanation:

During the preconfigure step of a PowerStore block import migration, network connectivity is set up. This step is crucial for ensuring that the source system can communicate with the PowerStore system, which is necessary for a successful data transfer during the migration process. Here are the detailed steps involved in this phase:

Assess Network Requirements: Evaluate the existing network infrastructure to ensure it meets the requirements for a PowerStore migration. This includes checking the bandwidth, latency, and network topology.

Configure Network Settings: Set up the network settings on both the source and destination systems. This involves configuring IP addresses, subnet masks, gateways, and any other network parameters

required for connectivity.

Verify Connectivity: Use network diagnostic tools to verify that the source system can communicate with the PowerStore system. This might involve ping tests, traceroutes, and other connectivity checks.

Secure the Network: Implement necessary security measures to protect the data during migration. This can include setting up firewalls, VPNs, or other security protocols to ensure data integrity and confidentiality.

Establish Data Paths: Create and configure data paths between the source system and the PowerStore system. This ensures that data can be transferred efficiently and without interruption during the migration process.

Test Network Performance: Conduct performance tests to ensure that the network can handle the data transfer load. This helps identify and mitigate any potential bottlenecks that could affect the migration.

Reference:

[Dell Technologies PowerStore Documentation](#)

[Dell Entry-Level and Mid-Range Storage Community](#)

[Dell Storage Products Overview](#)

By ensuring that network connectivity is properly set up, the preconfigure step lays the foundation for a smooth and successful migration of block data to the PowerStore system.

Question: 2

The CPU utilization in a Dell Unity XT system is currently at 68%. What is the level of CPU utilization?

- A. Extremely High Utilization
- B. High Utilization
- C. Low Utilization
- D. Normal Utilization

Answer: B

Explanation:

Understanding CPU Utilization Levels:

Dell Unity XT systems categorize CPU utilization to help administrators assess the performance and health of their storage systems.

Utilization levels typically include categories like low, normal, high, and extremely high to simplify performance monitoring.

CPU Utilization in Dell Unity XT:

According to Dell Unity XT documentation, CPU utilization percentages are used to determine the performance state of the system.

A CPU utilization of 68% is considered high but not critical. This is because the system is handling a significant load but still within operational parameters.

Operational Thresholds:

Low Utilization: Below 30%

Normal Utilization: 30% - 60%

High Utilization: 60% - 80%

Extremely High Utilization: Above 80%

Given these thresholds, 68% falls into the "High Utilization" category.

Dell Midrange Storage Reference:

Dell Technologies Partner Program documentation on midrange storage provides detailed metrics and thresholds for system performance. This information is crucial for administrators to make informed decisions about system capacity and performance tuning.

Additional resources on Dell's community forums and official product pages for entry-level and midrange storage also support these utilization thresholds.

By categorizing CPU utilization as "High" at 68%, administrators can take necessary actions to optimize system performance or plan for additional capacity if needed.

Reference:

[Dell Unity XT Overview](#)

[Dell Entry-Level and Mid-Range Community](#)

[Dell Storage Products](#)

Question: 3

What command is used to gather data from Brocade switches?

- A. show all
- B. show tech-support details
- C. supportshow

Answer: C

Explanation:

The command supportshow is used to gather comprehensive data from Brocade switches. This command collects extensive diagnostic information about the switch, including configuration details, hardware status, and various operational statistics. Here's a step-by-step explanation of the supportshow command:

Access the Brocade Switch: Connect to the Brocade switch via SSH or through a console connection.

Enter the Command: Type supportshow in the command line interface (CLI) and press Enter.

Data Collection: The command will begin collecting a wide array of information from the switch. This includes:

Hardware and firmware versions

Port status and configurations

Zoning information

Error logs and statistics

Environmental data (temperature, power supply status, etc.)

Output Review: The gathered data is typically extensive and can be used for troubleshooting, performance analysis, and ensuring the switch is operating correctly.

Save the Output: The output can be redirected to a file for later review or analysis by technical support teams.

Reference:

[Dell Technologies PowerStore Documentation](#)

[Dell Entry-Level and Mid-Range Storage Community](#)

[Dell Storage Products Overview](#)

Question: 4

When designing a system to support Oracle OLAP Workload for a new customer with no I/O details, what are the default selections for the I/O profile input in sizer?

- A. Sequential Read: 70% at 8 KiB Sequential Write: 30% at 8 KiB
- B. Random Read: 70% at 8 KiB Random Write: 30% at 8 KiB
- C. Sequential Read: 80% at 8 KiB Sequential Write: 20% at 8 KiB
- D. Random Read: 50% at 32 KiB Random Write: 50% at 32 KiB

Answer: B

Explanation:

When designing a system to support an Oracle OLAP (Online Analytical Processing) Workload for a new customer without specific I/O details, the default selections for the I/O profile input in the sizer tool are typically set to "Random Read: 70% at 8 KiB, Random Write: 30% at 8 KiB." Here's a detailed explanation:

Oracle OLAP Workload Characteristics: OLAP workloads generally involve heavy read operations with a focus on retrieving and analyzing large datasets. The read operations are often random due to the nature of querying various parts of the data.

Default I/O Profile Selection:

Random Read: 70% - This reflects the high read intensity typical of OLAP workloads, where data is accessed non-sequentially.

Random Write: 30% - Although less frequent than reads, writes still occur randomly as new data is added or existing data is updated.

Sizer Tool Configuration: The Dell sizer tool uses these default I/O profile settings to estimate the required storage performance and capacity. This helps in designing a system that can handle the expected workload efficiently.

Performance Optimization: The chosen I/O profile ensures that the storage system is optimized for the specific access patterns of OLAP workloads, providing a balanced approach to handling both reads and writes.

Reference:

[Dell Technologies PowerStore Documentation](#)

[Dell Entry-Level and Mid-Range Storage Community](#)

[Dell Storage Products Overview](#)

By understanding the default I/O profile settings used by the sizer tool, system designers can ensure that the storage solution is well-suited to meet the demands of Oracle OLAP workloads.

Question: 5

While sizing a PowerStore solution with 5 TB usable capacity and 5000 IOPS for a Test/Development Application using NVMe drives, a solution architect observes that the recommended solution is 10 times the requested capacity. What action must the architect take?

- A. Change the DRE Tolerance Parity in System Editor.
- B. Configure the Unity solution as an alternative.
- C. Change the DRR of the solution.

Answer: C

Explanation:

Understanding DRR (Data Reduction Ratio):

The Data Reduction Ratio (DRR) is a measure used in storage solutions to determine the efficiency of data reduction technologies such as deduplication and compression.

In NVMe-based systems like PowerStore, achieving an optimal DRR is crucial for balancing capacity and performance.

Sizing PowerStore Solutions:

When sizing a PowerStore solution, it's essential to align the provided capacity and performance (IOPS) with the customer's requirements.

If the recommended solution vastly exceeds the requested capacity (10 times more in this case), it indicates an overly conservative DRR assumption.

Recommended Action:

The architect should adjust the DRR to reflect a more accurate estimate that aligns with the customer's actual data reduction potential.

This adjustment ensures that the solution is cost-effective and meets the performance criteria without excessive over-provisioning.

Dell Midrange Storage Reference:

Dell Technologies documentation on PowerStore solutions provides detailed guidelines on DRR configurations and best practices.

Resources include sizing tools, system editor settings, and real-world case studies that illustrate appropriate DRR adjustments.

Reference:

[Dell PowerStore Overview](#)

[Dell Community on Midrange Storage](#)

[Dell Storage Product Information](#)